

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Architectures For And Methods Of Providing Network-based Software Extensions**

Inventor(s):

Paul Erickson  
Oliver Fisher

ATTORNEY'S DOCKET NO. ms1-586us

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25

WY  
O

## RELATED APPLICATIONS

The following patent applications are related to the present application, are assigned to the assignee of this patent application, and are expressly incorporated by reference herein:

- U.S. Patent Application Serial No. \_\_\_\_\_, entitled "Single Window Navigation Methods and Systems", bearing attorney docket number MS1-560us, and filed on the same date as this patent application;
- U.S. Patent Application Serial No. \_\_\_\_\_, entitled "Methods and Systems of Providing Information to Computer Users", bearing attorney docket number MS1-557us, and filed on the same date as this patent application;
- U.S. Patent Application Serial No. \_\_\_\_\_, entitled "Methods, Systems, Architectures and Data Structures For Delivering Software via a Network", bearing attorney docket number MS1-559us, and filed on the same date as this patent application;
- U.S. Patent Application Serial No. \_\_\_\_\_, entitled "Network-based Software Extensions", bearing attorney docket number MS1-563us, and filed on the same date as this patent application;
- U.S. Patent Application Serial No. \_\_\_\_\_, entitled "Authoring Arbitrary XML Documents Using DHTML and XSLT", bearing attorney docket number MS1-583us, and filed on the same date as this patent application;
- U.S. Patent Application Serial No. \_\_\_\_\_, entitled "Task Sensitive Methods And Systems For Displaying Command Sets", bearing attorney docket number MS1-562us, and filed on the same date as this patent application.

## TECHNICAL FIELD

This invention relates to methods and systems for providing software via a network. More specifically, the invention pertains to Internet-based delivery of software.

## BACKGROUND

Installation of traditional PC applications requires physical media, such as a disk or CD-ROM that must be physically inserted into a computer in order for software to be loaded onto a user's computer. Typically, this process requires the user to enter settings information that can be confusing to the user. Once the software is installed, it is typically fixed in terms of its location and functionality. When the software is updated, the user must typically purchase additional physical media and repeat the installation process so that they can use the updated software. In this model, the software is fixed in its association with the computer on which it was installed. If a user moves to another computer, they will not be able to use the specific software on their machine without repeating the installation process.

As computing continues to evolve into the environment of computer networks such as the Internet, it has become clear that the traditional software delivery model described above is inadequate to meet the demands of consumers who desire dynamic, flexible, and adaptable software on-demand. Network-based software delivery is becoming the subject of increasing focus by those who develop and deliver software. Unlocking the potential for network-based software delivery will require smart, innovative and streamlined solutions, especially in situations where bandwidth may be limited.

Accordingly, this invention arose out of concerns associated with providing new software delivery models that are particularly well-suited for network-based software delivery, e.g. delivery via the Internet.

1 **SUMMARY**

2 Methods and systems for network-based software delivery are described.  
3 In one embodiment, an application program or software platform resides on a  
4 client. The program or platform is configured so that it is extensible based on  
5 software extensions that are deliverable over a network such as the Internet.  
6 Various extensions can be developed by third party developers for incorporation  
7 into the program or platform.

8 In one described embodiment, extension files that comprise a software  
9 extension are hosted on a network server such as an Internet server. These include  
10 descriptor files that describe aspects of the software extension. These descriptor  
11 files logically describe an extension to the program or platform and designate the  
12 location of other extension files. Extensions are incorporated on a client by  
13 navigating to a particular network or Internet site through which the extensions  
14 can be accessed. The files describing the extension files are downloaded on the  
15 client. These files tell the client where, when and how the particular extension can  
16 be plugged into the program or platform, as well as where to find the appropriate  
17 extension files and how to download them. The extension files are then  
18 downloaded and incorporated into the platform.

19 In one embodiment, an inventive software architecture is provided for  
20 handling and consolidating particular types of descriptive files associated with  
21 various extensions. A filtering mechanism called attachment points, is used to  
22 create handlers for the different descriptive types of files that define a software  
23 extension. Each of these handlers is known as an attachment manager.  
24 Attachment managers are provided for each extensible feature type. The  
25 attachment managers interpret data from extensions files which are supplied by

1 attachment points. In addition to predefined attachment managers, custom  
2 attachment managers can be created using data from attachment points. When an  
3 extension extends a particular feature type, the attachment points ensure that only  
4 the appropriate attachment manager is notified so that the feature type can be  
5 incorporated into the program or platform efficiently.

6

7 **BRIEF DESCRIPTION OF THE DRAWINGS**

8 Fig. 1 is a high level view of a system that can be utilized in accordance  
9 with one described embodiment.

10 Fig. 2 is an exemplary computer system that can be utilized in accordance  
11 with the described embodiment.

12 Fig. 3 is a diagram of an exemplary EDF and PKG in accordance with one  
13 described embodiment.

14 Fig. 4 shows a portion of an EDF in accordance with the described  
15 embodiment.

16 Fig. 5 shows a portion of an EDF schema in accordance with the described  
17 embodiment.

18 Fig. 6 shows a portion of a PKG in accordance with the described  
19 embodiment.

20 Fig. 7 is a block diagram illustrating how file hashes can be used for  
21 versioning in accordance with one embodiment.

22 Fig. 8 is a block diagram that illustrates two exemplary package objects in  
23 accordance with one embodiment.

24 Fig. 9 is a flow diagram that describes steps in a method in accordance with  
25 one described embodiment.

1 Fig. 10 is a flow diagram that describes steps in a method in accordance  
2 with one described embodiment.

3 Fig. 11 is a block diagram that illustrates an exemplary package manifest  
4 creation tool in accordance with one described embodiment.

5 Fig. 12 is a flow diagram that describes steps in a method in accordance  
6 with one described embodiment.

7 Fig. 13 is a flow diagram that describes steps in a method in accordance  
8 with one described embodiment.

9 Fig. 14 is a flow diagram of steps in a method in accordance with the  
10 described embodiment.

11 Fig. 15 is a flow diagram of steps in a method in accordance with the  
12 described embodiment.

13 Fig. 16 shows a portion of an exemplary catalog structure in accordance  
14 with the described embodiment.

15 Fig. 17 is a block diagram of a software architecture in accordance with the  
16 described embodiment.

17 Fig. 18 is a flow diagram of steps in a method in accordance with the  
18 described embodiment.

19 Fig. 19 is a diagram that illustrates one aspect of attachment point  
20 architecture in accordance with the described embodiment.

21 Fig. 20 is a diagram that illustrates one aspect of the Fig. 17 architecture.

1      **DETAILED DESCRIPTION**

2      **Overview**

3      The methods and systems described just below provide a mechanism by  
4      which functionality can be added dynamically to an application program or  
5      software platform. Functionalities or “extensions” as they will be referred to, can  
6      be advantageously added via a network such as the Internet. Extensions, that can  
7      implement new features or add to existing features, can be added using only a  
8      network address, such as a URL, as a basis for extension installation. That is, all  
9      of the files that comprise an extension can be maintained on the network and  
10     accessed via one or more network sites.

11     Extensions can be described in a variety of ways. One way utilizes a  
12     hierarchical tag-based language which facilitates handling and use of the various  
13     extensions. In one particular implementation, a software platform is provided that  
14     can incorporate various functionalities. The software platform and the inventive  
15     architecture described below enable third and fourth party developers to develop  
16     extensions for the platform that can be easily and seamlessly incorporated into the  
17     platform without having any knowledge of (or relationship with) a hosting service.  
18     A third party developer is a developer who develops an extension for the platform.  
19     A fourth party developer might be a developer who develops an extension to a  
20     third party developer’s extension. Thus, the incorporation of third and fourth party  
21     extensions is essentially a transparent process, as far as developers are concerned.

22     Consider for example, Fig. 1 which shows a user’s computer 100 and  
23     several so-called extension sources 102, 104, and 106. The extension sources can  
24     comprise any entity from which a software extension can be obtained via a  
25     network. In an exemplary implementation, the network can comprise the Internet,

1 although other networks (e.g. LANs and WANs) can certainly be utilized.  
2 Extension sources can include, without limitation, business entities such as retail  
3 stores that might maintain a network site. In one implementation, a user can  
4 execute software on their computer that provides an application program or  
5 software platform. In this document, the terms “application program” and  
6 “software platform” will be used interchangeably. Each of the different extension  
7 sources 102-106 can provide software extensions that can plug into the software  
8 platform that is executing on the user’s machine. These extensions are deliverable  
9 via a network such as the Internet, and assist in providing applications that can be  
10 executed on the user’s machine. In the described embodiment, the extensions are  
11 logically described in XML which is in line with emerging industry standards.  
12 Additionally, the use of XML assists in the future discoverability of extensions by  
13 promoting XML DOM properties on the Internet. It will be appreciated, however,  
14 that any suitable format can be used for describing the extensions, e.g. a binary  
15 description could be used.

16 Extensions can be delivered from any number of different extension  
17 sources. The inventive methods and systems provide a streamlined and organized  
18 way to handle the provided extensions. The use of XML advantageously enables  
19 efficient handling of extensions from multiple different extension sources, without  
20 unduly taxing the software components that utilize specific portions of an  
21 extension or extensions.

22 In one particular implementation, the software platform on the user’s  
23 machine provides various different integrated functionalities that enable a user to  
24 accomplish different document-centric tasks. An exemplary system is described in  
25

1 the U.S. Patent Application entitled "Single Window Navigation Methods and  
2 Systems", incorporated by reference above.

3

4 **Exemplary Computer Environment**

5 The embodiment described just below can be employed in connection with  
6 various computer systems. A computer system, for purposes of this document,  
7 can be considered as any computing device that includes some type of processor,  
8 i.e. a microprocessor, and some type of operating system. Thus, a computer  
9 system can be construed to include, without limitation, traditional desktop  
10 computers, more powerful servers, various hand-held devices such as cell phones,  
11 pocket-held computer devices and the like.

12 Fig. 2 shows but one exemplary computer system that can be used to  
13 implement the embodiments described herein. Computer 130 includes one or  
14 more processors or processing units 132, a system memory 134, and a bus 136 that  
15 couples various system components including the system memory 134 to  
16 processors 132. The bus 136 represents one or more of any of several types of bus  
17 structures, including a memory bus or memory controller, a peripheral bus, an  
18 accelerated graphics port, and a processor or local bus using any of a variety of  
19 bus architectures. The system memory 134 includes read only memory (ROM)  
20 138 and random access memory (RAM) 140. A basic input/output system (BIOS)  
21 142, containing the basic routines that help to transfer information between  
22 elements within computer 130, such as during start-up, is stored in ROM 138.

23 Computer 130 further includes a hard disk drive 144 for reading from and  
24 writing to a hard disk (not shown), a magnetic disk drive 146 for reading from and  
25 writing to a removable magnetic disk 148, and an optical disk drive 150 for

1 reading from or writing to a removable optical disk 152 such as a CD ROM or  
2 other optical media. The hard disk drive 144, magnetic disk drive 146, and optical  
3 disk drive 150 are connected to the bus 136 by an SCSI interface 154 or some  
4 other appropriate interface. The drives and their associated computer-readable  
5 media provide nonvolatile storage of computer-readable instructions, data  
6 structures, program modules and other data for computer 130. Although the  
7 exemplary environment described herein employs a hard disk, a removable  
8 magnetic disk 148 and a removable optical disk 152, it should be appreciated by  
9 those skilled in the art that other types of computer-readable media which can  
10 store data that is accessible by a computer, such as magnetic cassettes, flash  
11 memory cards, digital video disks, random access memories (RAMs), read only  
12 memories (ROMs), and the like, may also be used in the exemplary operating  
13 environment.

14 A number of program modules may be stored on the hard disk 144,  
15 magnetic disk 148, optical disk 152, ROM 138, or RAM 140, including an  
16 operating system 158, one or more application programs 160, other program  
17 modules 162, and program data 164. A user may enter commands and  
18 information into computer 130 through input devices such as a keyboard 166 and a  
19 pointing device 168. Other input devices (not shown) may include a microphone,  
20 joystick, game pad, satellite dish, scanner, or the like. These and other input  
21 devices are connected to the processing unit 132 through an interface 170 that is  
22 coupled to the bus 136. A monitor 172 or other type of display device is also  
23 connected to the bus 136 via an interface, such as a video adapter 174. In addition  
24 to the monitor, personal computers typically include other peripheral output  
25 devices (not shown) such as speakers and printers.

1 Computer 130 commonly operates in a networked environment using  
2 logical connections to one or more remote computers, such as a remote computer  
3 176. The remote computer 176 may be another personal computer, a server, a  
4 router, a network PC, a peer device or other common network node, and typically  
5 includes many or all of the elements described above relative to computer 130,  
6 although only a memory storage device 178 has been illustrated in Fig. 2. The  
7 logical connections depicted in Fig. 2 include a local area network (LAN) 180 and  
8 a wide area network (WAN) 182. Such networking environments are  
9 commonplace in offices, enterprise-wide computer networks, intranets, and the  
10 Internet.

11 When used in a LAN networking environment, computer 130 is connected  
12 to the local network 180 through a network interface or adapter 184. When used  
13 in a WAN networking environment, computer 130 typically includes a modem 186  
14 or other means for establishing communications over the wide area network 182,  
15 such as the Internet. The modem 186, which may be internal or external, is  
16 connected to the bus 136 via a serial port interface 156. In a networked  
17 environment, program modules depicted relative to the personal computer 130, or  
18 portions thereof, may be stored in the remote memory storage device. It will be  
19 appreciated that the network connections shown are exemplary and other means of  
20 establishing a communications link between the computers may be used.

21 Generally, the data processors of computer 130 are programmed by means  
22 of instructions stored at different times in the various computer-readable storage  
23 media of the computer. Programs and operating systems are typically distributed,  
24 for example, on floppy disks or CD-ROMs. From there, they are installed or  
25 loaded into the secondary memory of a computer. At execution, they are loaded at

1 least partially into the computer's primary electronic memory. The invention  
2 described herein includes these and other various types of computer-readable  
3 storage media when such media contain instructions or programs for implementing  
4 the steps described below in conjunction with a microprocessor or other data  
5 processor. The invention also includes the computer itself when programmed  
6 according to the methods and techniques described below.

7 For purposes of illustration, programs and other executable program  
8 components such as the operating system are illustrated herein as discrete blocks,  
9 although it is recognized that such programs and components reside at various  
10 times in different storage components of the computer, and are executed by the  
11 data processor(s) of the computer.

001230475000

1                    **Extensions**

2                    An “extension”, as used in this document, will be considered to include,  
3 without limitation, software functionality and content that can be added to an  
4 application program or software platform. These additions typically provide some  
5 type of functionality that the application program may not have had before the  
6 extension was incorporated, or alter the behavior of at least one existing feature.  
7 The extension is incorporated or integrated directly into the application program in  
8 a way that changes, to some degree, the manner in which the application program  
9 behaves or operates. Extensions provide dynamically added content and can  
10 provide applications (such as an email application), plug-ins to extend existing  
11 applications (like a fax plug-in to an email application), or simply web pages, to  
12 name just a few.

13                    In the described embodiment, extensions are described using XML, an  
14 industry-standard, text-based markup language. XML greatly facilitates the  
15 extensibility of software content. Specifically, various extensions can be authored  
16 by third parties and described in XML in a manner that enables the extensions to  
17 be readily integrated into application programs. It should be appreciated, however,  
18 the XML constitutes but one exemplary way of describing and using the  
19 extensions. Other ways can, of course, be used.

20                    **Exemplary Extension Organization**

21                    In the described embodiment, extensions are organized in three separate but  
22 related portions: an Extension Definition File (EDF), a Package Manifest (PKG),  
23 and the code, components, or “bits” that make up or define the extension. An EDF  
24 can be, but need not be associated with a URL (Universal Resource Locator) that

1 provides a way for a client to access the EDF. By convention and choice, the PKG  
2 file is located at the same URL as the EDF. It will be appreciated that the  
3 described EDFs and PKG are each not required for the other to be used. It just so  
4 happens that, in the example that is given in this document, the two are employed  
5 together. To that end, each of these features can be separately and independently  
6 employed.

7 EDFs describe logical attachments to an application program or software  
8 platform, while PKGs specify the physical files and resources that are used in an  
9 extension. There can be a one to one correspondence between EDFs and PKGs.

10 Fig. 3 shows an exemplary organization 300 that includes an EDF 302 and  
11 a corresponding package manifest (PKG) 304. In the illustrated example, the EDF  
12 302 uses XML to describe the logical attachments or extensions to an application  
13 program. The corresponding PKG 304 specifies the physical files and resources  
14 that are associated with a particular extension. Exemplary file types are shown to  
15 the right of PKG 304 and include, without limitation, HTM, GIF, UDO, XML,  
16 DLL, and various other types of files.

## 17 18 **Extension Definition File (EDF)**

19 In the described example, an EDF is an XML file that logically describes an  
20 extension. For example, the EDF can describe HTML that makes up a user  
21 interface (UI), the objects that contain code for implementing various functions,  
22 and the like. The EDF can also contain all or part of the functionality that  
23 comprises an extension. For instance, the HTML that describes a toolbar could be  
24 incorporated directly into an EDF file, and a toolbar attachment manager could  
25 read it from the EDF file, instead of from a URL. The information contained in

1 the EDF is processed and (together with the information contained in the PKG),  
2 the appropriate files are automatically installed on a user's computer. This is done  
3 unobtrusively without manipulating the computer's persisted settings, as might be  
4 found in the user's system registry.

5 An EDF, implemented in XML, contains various tags that are associated  
6 with various extensions. The various tags can correspond to:

- 7 • User interface elements
- 8 • Behaviors/Components/Objects
- 9 • Store Elements
- 10 • User-defined objects
- 11 • Or anything else that represents a point of extensibility in the application or platform

12 EDFs advantageously have an "open schema" which means that third party  
13 developers can extend the extension mechanism and include their own extensions  
14 by creating their own tags. Additionally, extensions can themselves be extended  
15 by other developers. EDFs can also have one or more predefined tags. Exemplary  
16 predefined XML tags for user interface elements can include tags for feature types  
17 such as: tool bars, accelerators, menu items, and themes. These feature types are  
18 utilized in the single navigable window application incorporated by reference  
19 above and defined in the table immediately below:

1 <b>Feature Type</b>	2 <b>Definition</b>
2 Tool Bars	3 Horizontal command containers above the document area.
3 Accelerators	4 Keyboard shortcuts for commands
4 Menu Items	5 Pop-up or drop-down menu choices that third parties can add to well-known, named menu attachments in the platform
5 Themes	6 A data-driven way to provide overrides for well-known resources of the platform, such as default buttons or default style sheet

7  
Table 1

8  
9  
Exemplary predefined XML tags for behaviors/components/objects include  
10 tags for Services. These feature types are utilized in the single navigable window  
11 application incorporated by reference above and defined in the table immediately  
12 below:

14 <b>Feature Type</b>	15 <b>Definition</b>
15 Services	16 Services are objects that extend existing objects (such as 17 the application, window, or document) in the application 18 or platform's Object Model. For example, editing 19 functions use Object Model Attachments attached to the window or document that maintain document context and editing state per-window. These can also include Object Model Attachments attached to the application (such as a spellchecker dictionary object)

20  
Table 2

21  
22  
Exemplary predefined XML tags for store elements include tags for content  
23 classes and offline data sources. These feature types are utilized in the single  
24 navigable window application incorporated by reference above and defined in the  
25 table immediately below:

Feature Type	Definition
Content Classes	Allow extension writers to define new types of XML documents with new schemas.
Offline Data Sources	Allow for extension writers to define store replication instructions in an EDF.

Table 3

### EDF Schema

In the described embodiment, the EDFs have a particular XML schema that is utilized. The schema comprises collections of XML tags that are arranged in a hierarchical organization to facilitate information dissemination to software components that need certain extensions. In the described embodiment, the outer (encompassing tag) for EDFs is an “extension” tag.

Fig. 4 shows an exemplary extension tag. “extension” tags can include one or more of the following attributes, all of which are optional:

Attribute	Definition
urn	ID for the extension. It allows extension writers to specify relative locations for content in EDFs without using relative paths or fixed URLs. It also allows hosting administrators to move around extensions on servers without breaking any links.
name	Name that can be used in a status bar or message display
version	Vendor-determined version number for the extension.
lastUpdate	Date/time that the EDF was last modified.
description	Brief description of the extension.

Table 4

Within the “extension” outer tag are one or more child tags, also referred to as “top level tags”. These top level tags are each associated with a feature type

1 that can be added by a particular extension. Exemplary feature types are discussed  
2 in connection with Tables 1-3 above. Underneath each top level tag there can be  
3 one or more child tags that are individually associated with a particular feature of  
4 the feature type that is to be added by a particular extension.

5 Fig. 5 shows an exemplary XML schema organization in accordance with  
6 this embodiment. For each top level tag in an EDF, there is an associated  
7 attachment manager which is a software component that receives data associated  
8 with the tag so that the data can be used to incorporate the extension into the  
9 platform or application program. Different attachment managers may interpret the  
10 data from the tag in different ways to provide different types of extensibility, so  
11 different top level tags will contain different types of data in different structures.  
12 This will become more evident in the “Architecture” section below. Note that the  
13 “**edf:**” XML namespace qualifier enables support of an open schema where  
14 extensions can provide their own tags and corresponding attachment managers.  
15 Tags within the **edf** namespace are used by built-in attachment managers in the  
16 application or software platform. Tags in other namespaces are used by third-  
17 parties to provide additional points of extensibility.

18

### 19 **Package Manifest (PKG file)**

20 The package manifests (PKGs) assist in organizing the downloading of  
21 software in the form of multiple files over a network such as the Internet. The  
22 PKGs are advantageously employed, in the example given in this document, with  
23 EDFs. As pointed out above, however, the techniques discussed in connection  
24 with the PKGs can be deployed independently of EDFs and in any suitable  
25 scenario where it is desirable to deliver software over a network such as the

1 Internet. While the EDFs describe the logical attachment of extensions into an  
2 application program or platform, the package manifest's role is to assist in one or  
3 more of: organized delivery, validation and/or, updating of files associated with the  
4 various extensions that can be provided.

5 In designing a delivery mechanism for Web-assisted delivery of software  
6 content or files, several considerations are of interest.

7 Whenever possible, it is desirable to reduce the size of required downloads  
8 during update and installation operations. To address this consideration, software  
9 content is broken into multiple packages. Each package contains a group of one or  
10 more files that implement a common or well-defined functionality. By breaking  
11 the content into individual packages, the size of the required download during  
12 installation and update can be minimized. Each package is then described by a  
13 package manifest (PKG) that includes file information such as file locations and  
14 hash values that can be used for validation or security and versioning.

15 It is also desirable to give the end user a Web-like experience. To do this,  
16 extensions are loaded in a manner that feels to a user more like they are loading a  
17 web page, rather than traditional software packages where the user has to wait  
18 until the entire package is loaded before they can interact with it. In the described  
19 embodiment, users are given a web-like experience by streaming the extension  
20 files down to the client so that a user can begin to interact with an application  
21 program much sooner than if they had to wait for the entire software application  
22 program to load. For example, if there are user interface (UI) image files  
23 streaming down, the user can see the UI as the files stream in. Consider, for  
24 example, the single application program having the multiple different  
25 functionalities that is described in the patent application incorporated by reference

1 above. A user might browse to an email functionality and download the files that  
2 are necessary to interact with the email functionality. Files that are associated  
3 with another different functionality would then be downloaded after the files  
4 associated with the email functionality. In this way, the user can begin to operate  
5 within a particular functionality without having to wait for all of the files  
6 associated with all of the other functionalities.

7 Another consideration of interest pertains to the efficiency with which the  
8 extension files or “bits” are delivered to the client. To address this consideration,  
9 the described embodiment utilizes a couple of different download approaches: a  
10 throttled download and a background download. Throttled downloading conducts  
11 download operations while taking into account the available bandwidth and type  
12 of media over which the files are being transferred. Any suitable throttled  
13 download process can be employed and will be understood by those of skill in the  
14 art. Background download is conducted while the user is working within the  
15 application program and is implemented by allocating a background thread so that  
16 the user can continue their work. One optimization that is utilized is that packages  
17 are prioritized and delivered in accordance with what a user might be working on.

18 Another consideration is associated with optimizing the user’s computing  
19 experience. Here, the user’s experience is optimized by making available the most  
20 common scenarios for the user. This is effectively directed to giving the user the  
21 functionality that they want first, and then, through the background download  
22 process, providing the code that implements functionalities that the user might use  
23 in the future. To determine which functionalities a user desires to have first, an  
24 automated scenario-based packaging process is provided that runs against file  
25 usage logs from scripted scenarios.

1 All of these considerations and the inventive solutions directed to  
2 addressing the considerations are discussed in detail in the sections that follow  
3 below.

5 **Package Manifest Definition**

6 In the described embodiment, a package manifest (PKG file) comprises a  
7 list of files that are utilized in a package. The list is advantageously compressed  
8 somewhat and digitally signed. Each package manifest can contain a list of one or  
9 more files each of which can include an associated hash, as well as download  
10 directives that control caching of the files. Once an extension is authored, a  
11 software tool can be used to generate the package manifest.

12 In addition, the package manifest can specify several other pieces of  
13 information:

15 • FILE GROUP

16 All files in an extension can be labeled according to a number of predefined  
17 file groups. The file group of a particular file determines when the particular file  
18 gets downloaded, where it is stored on the client, and how it gets packaged. In the  
19 described embodiment, four predefined file groups are provided and are listed and  
20 described in the table immediately below:

Group name	When downloaded	Where stored on the client	Packaging	Content
Required	Downloaded before any other files in the extension.	NetDocs package cache	All required files in an extension are	DLLs included so that a user will not have to wait for a prolonged

			packaged together as a CAB* file.	period of time before clicking on a UI element
1	Offline	Offline files start getting downloaded as soon as Required are down. Providing the user stays on line long enough, these files will all get downloaded and will later be available for offline use.	NetDocs package cache	File are sent down individually.
2	On demand	Only downloaded when they are requested for the first time.	NetDocs package cache	Files are sent down individually.
3	Online only	Downloaded on demand. Content is only available when the user is online.	WinInet Cache	Content that is not to be provided offline. Examples include help pages and other content that can consume a large amount of disk space.

11 \* CAB stands for the CABinet technology that Internet Explorer uses to package bits for download. CAB files average from 2 to 3:1 compression, and are optimized for quick expansion. CAB files have the added security benefit that they are easily signed.

13           •     FILE DOWNLOAD PRIORITY

14           Files in each group are listed according to the order in which they should be  
 15 downloaded. This download order is implicit in the ordering of the files in the  
 16 package manifest, an example of which is shown in Fig. 6.

18           •     HASH VALUE FOR SECURITY/VERSIONING

19           Individual files in the package manifest can have an associated hash value.  
 20 Each hash value is generated by running the file through an encryption algorithm.  
 21 An exemplary encryption algorithm is Microsoft's CryptoAPI. In the illustrated  
 22 example, each file can be listed with a base 64-encoded hash value, so that the file  
 23 can be validated once the content arrives at the client. Specifically the package  
 24 manifest is sent to the client in a secure manner (i.e. it is digitally signed). The  
 25

1 package manifest contains the hash values for individual files. When the  
2 individual files are received by the client, each of the files can be run through the  
3 same Crypto API that was used to provide the hash values in the package  
4 manifest. If the hash values for a particular file compare favorably, then the file  
5 has not been altered and is secure.

6 When a file is updated, hash values can serve a useful purpose in  
7 identifying files that have not been changed between different versions of an  
8 extension. Consider Fig. 7, for example. There, an old directory 700 in a client  
9 package cache contains package A which include two files—file 1 with hash = x,  
10 and file 2 with hash = y. Assume that this package is associated with an older  
11 version of an extension. When an updated version is produced, its package  
12 manifest is delivered to the client. The updated extension version is represented  
13 on a source directory of a code or web server 704. The package manifest includes  
14 the hash values for all of the files in the new extension version. A new client  
15 destination directory 702 is defined for all of the files of the new extension. If any  
16 of the hash values for files in the new extension are the same as the hash values of  
17 the files in the old directory 700, then those files can be copied directly from the  
18 old directory 700 to the new destination directory 702. In this example, file 1's  
19 hash value is the same as the hash value for file 1 on the source directory 704, so it  
20 can be copied into the new destination directory 702. File 2's hash value, however  
21 is different from the hash value for file 2 on the source directory, so it is not copied  
22 from the old directory 700. Rather, file 2 is downloaded from the code server. A  
23 new file 3 has been added and is also downloaded from the code server. Hence, in  
24 this example, a new version of an extension resulted in a download of less than all  
25 of the files in the extension version. This is because hash values for each of the

1 files in the old extension version were able to be compared with hash values of the  
2 files in the new extension version. Those hash values that are the same indicate  
3 files that have not changed as between versions.

4 Using hash values for versioning has two important advantages over  
5 traditional versioning schemes. First, the update process is automatic. That is,  
6 with an explicit version number, it is possible to forget to update the version  
7 number when shipping a new release of a file. Using hash values avoids this  
8 problem. Second, versioning does not rely on file types. Specifically, traditional  
9 versioning schemes commonly embed version information within files; however,  
10 not all files (e.g. GIF files) support embedded version information. In the present  
11 example, using hash values for versioning does not depend on whether a particular  
12 file type supports or does not support embedded version information. In addition,  
13 the version information can be stored separately from the file itself. Thus, access  
14 to actual file to determine whether it is current is not needed.

15

- 16 • TOTAL STORAGE SIZE OF PACKAGE

17 The total storage size of a package is useful at download time to verify that  
18 the user has sufficient disk space.

19

- 20 • ClassID's FOR THE DLLs

21 Listing ClassIDs for each DLL is necessary to enable script writers to  
22 create classes by scripting against the OM. Additionally, this enables a  
23 determination of which package contains the code for a particular class.

## 1 DLL LOAD DEPENDENCIES

2 The reason for the dependencies section is to allow for legacy code that  
3 relies on being loaded by virtue of being in the search path of some other dll. In  
4 this case we have to make sure that the dependency dll is in the package cache  
5 directory before the dependant dll is loaded. Fig. 6 shows an exemplary package  
6 manifest 600 that is defined in a hierarchical tag-based language. Advantageously,  
7 the tag-based language comprises XML which is desirably extensible and flexible.  
8 In this example, a number of tags are provided in a hierarchical arrangement. The  
9 “package” tag contains information about the size of the package. The “files” tag  
10 is a child of the “package” tag and contains information about the file groups that  
11 are contained in that particular package. The “file” tag is a child of the “group”  
12 tag and contains information about specific files that comprise the extension, i.e.  
13 file name and hash value. A “dependency” tag is provided as a child of the “file”  
14 tag and lists any dependencies as discussed above. A “COMClass” tag is also  
15 provided as a child of the “file” tag and contains IDs as mentioned above. The  
16 ordering of the file groups in this schema implicitly defines the download order of  
17 the files.

## 18

### 19 Package Delivery

20 To optimize package delivery, two different delivery schemes are utilized.  
21 First, a throttled download approach is utilized using known throttling download  
22 techniques. Here, considerations such as available bandwidth and media over  
23 which the extensions are being provided are considered.

24 Second, a background download approach is utilized. Background  
25 downloads enable a user to continue to work within an application program while

1 content is downloaded. Foreground downloads are used when the user has  
2 explicitly requested a file/extension by clicking, for example, on an extension link,  
3 or requested an action, for example, by clicking on the “Compose” mail button,  
4 that requires download of files which are not available locally.

5 Along with background downloads, a queue management feature is  
6 provided. Specifically, when an extension must be installed or updated, a package  
7 manager, which is essentially a software component that manages packages, is  
8 provided with the following information:

- 9 • URL for the package manifest information on a code server
- 10 • URN for package destination directory in the package cache at the client
- 11 • (Optional) URN for the old package directory (if one exists) in the  
package cache

12 From this information, the package manager creates a package object and  
13 adds the package object to a download queue. The download queue is designed  
14 for easy rearrangement of a package download order. Consider, for example, Fig.  
15 8 which shows a portion of a download queue 800 that contains two package  
16 objects—package object 802 (corresponding to package A) and package object  
17 804 (corresponding to package B). The package objects maintain a list of which  
18 files of a corresponding package have been downloaded or installed. In the  
19 present example, files 1 and 2 from package A have been installed while file 3 has  
20 not been installed; and files 1, 2, and 3 have not been installed from package B.  
21 The download queue can be rearranged based on what the user is doing. That is,  
22 based on the actions that a user takes, the priority of files that are to be  
23 downloaded can be changed. In this example, the package manager is designed to  
24 process the first uninstalled file in the package at the head of the download queue.  
25

1 If, however, the user starts to use a file in an extension that is different from the  
2 extension whose files are at the head of the download queue, the corresponding  
3 package for the file that the user has started to use can be moved to the head of the  
4 download queue. Because a file's package is specified by its URN, the file's  
5 package can be quickly identified and located in the download queue. For  
6 example, and considering Fig. 8, if one of the files in package B is requested  
7 before the package manager has started to install package A's third file, then  
8 package B will be moved to the head of the download queue.

9 Fig. 9 is a flow diagram that describes steps in a download queue  
10 management method in accordance with the described example. The method can  
11 be implemented in any suitable hardware, software, firmware or combination  
12 thereof. In the present example, the method is implemented in software.

13 Step 900 receives one or more requests for an extension. The requests can  
14 be generated in any suitable way. Step 902 creates a package object that  
15 corresponds to each extension package that is to be downloaded. Step 904  
16 arranges the package objects in a download queue. Step 906 then downloads files  
17 corresponding to the package objects in the download queue. This step can be  
18 implemented by, for example, starting at the head of the download queue and  
19 downloading files until all of the files for a package object have been downloaded,  
20 and then moving to the next package object. Step 908 ascertains whether a user  
21 action requires a file that is not described in the current package object. If the  
22 user's action does not require a file not described by the current package object,  
23 then the method branches back to step 906 and continues to download files  
24 associated with the current package object. If, on the other hand, the user's action  
25 requires a files that is not described in the current package object, then step 910

1 moves the package object associated with the required file to the head of the  
2 download queue and begins to download files associated with this newly-  
3 repositioned package object. This step can be implemented by ascertaining which  
4 package object is associated with the required file by ascertaining the URN  
5 associated with the file. This URN specifies the file's package so that its package  
6 object can be quickly located and moved to the front of the download queue.

7

8 Package Creation

9 One of the innovative features of the described embodiment is its  
10 extensibility. That is, a software platform is provided in the form of an application  
11 program that can be extended by various third-party user-defined extensions.  
12 These extensions are delivered via the Web and are integrated directly into the  
13 software platform. In order to provide an organized delivery process, packages  
14 should be created in a uniform manner so that they can be predictably handled and  
15 integrated into the software platform.

16 In accordance with the described embodiment, each package should  
17 correspond to an end-user feature. For example, in the patent application  
18 incorporated by reference above, separate packages are provided for each of the  
19 email, contacts, document authoring, and planner functionalities. If packages that  
20 do not depend on one another share a dependency, then this shared dependency  
21 should become its own package. For example, there is no reason why the email  
22 and document authoring functionalities should depend on one another, yet both of  
23 them require the ability to publish content. By separating the publishing  
24 functionality into its own package, a certain amount of download order flexibility  
25 is preserved. Depending on what the user starts to work on, the files

1 corresponding to the email functionality or the document authoring can be  
2 downloaded first.

3 Fig. 10 is a flow diagram that describes steps in a package creation method  
4 in accordance with the described example. The method can be implemented in  
5 any suitable hardware, software, firmware or combination thereof. Portions of the  
6 method might, however, be implemented manually.

7 Step 1000 identifies end user features that are to be provided as extensions.  
8 Step 1002 identifies any shared dependencies among the end user features. Step  
9 1004 creates individual packages for the end user features. Step 1006 creates  
10 individual packages for any shared dependencies among the end user features.

11

12 Automated Package Manifest Creation Tool

13 Advantageously, and in accordance with one implementation, an automated  
14 package manifest tool is provided and takes various input parameters and  
15 automatically creates a package manifest. The tool can be available to third  
16 parties to assist them in creating a package manifest.

17 Fig. 11 shows an exemplary package manifest creation tool 1100 that is  
18 desirably implemented in software. In this specific example, the tool can take the  
19 following input parameters (some of which are optional):

20

21

22

- Extension directory
- File group information and DLL load dependencies (Optional)
- File usage statistics from scenario runs (Optional)

23 The extension directory input parameter specifies the directory containing  
24 all of the files that will be described by the package manifest. If this is the only  
25

1 parameter, then tool 1100 will generate a manifest in which the EDF and DLLs in  
2 the directory are listed in the “Required” set, and all other content is “Offline”.

3 The file group information and load dependencies parameter is optional. If  
4 an extension author has an idea of the categories in which his or her files should be  
5 placed, the categories should be specified here. For example, the author of the  
6 template manifest shown below knows that he wants his error handling GIF to be  
7 included in the package’s required set. His choices here will always be respected  
8 in the final manifest. Additionally, if the extension author knows of any DLL load  
9 dependencies, these should be specified here as well.

```
10 <?xml version="1.0"?>
11 <Package>
12   <Files>
13     <Group Name="required">
14       <File Name = "bar.dll"/>
15       <File Name = "foo.dll"/>
16       <Dependencies>
17         <File Name = "bar.dll"/>
18       </Dependencies>
19     </File>
20     <File Name = "myextension.edf"/>
21     <File Name = "errorhandling.gif">
22   </Group> ...
23 </Files>
24 </Package>
```

25 The file usage statistics from scenario runs parameter is an optional  
26 parameter. This parameter enables the file download priority to be determined  
27 based on scenario runs. A scenario is a script of tasks that the average user  
28 typically follows when using a product during a particular portion of product use.  
29 For example, one scenario might pertain to the tasks involved in sending an email

1 message (i.e. click “new mail” button, type in “TO” well, type in “Subject” well,  
2 etc.). In the described embodiment, file usage statistics from scenario runs are  
3 collected from running IIS logs on various scenarios. The different scenarios are  
4 directed to ensuring, with some degree of probabilistic support, that the file  
5 download order reflects, in some way, the files that will likely be used by the user  
6 first.

7 It will be appreciated that the file usage statistics can be provided  
8 dynamically by building a knowledge base that describes the actual tasks that  
9 people typically accomplish. The information that is maintained in the knowledge  
10 base can then be used to generate and adapt download scenarios that actually  
11 conform to patterns that are established across a user base.

12 If extension writers want to pass the package manifest creation tool 1100  
13 this information, they need to specify the log directory, as well as the start and end  
14 dates of the section of the log that the tool should analyze. For third parties, the  
15 download priority order within a group will be the order in which the group’s files  
16 were requested in the logs across all scenarios.

17 In one implementation, the approach is somewhat more sophisticated.  
18 Additional information (in addition to the scripted steps) is stored in the IIS logs  
19 and includes scenario priority and checkpoints. The scenario priority is a priority  
20 that is assigned for each scenario. So, for example, if one scenario is ten times  
21 more important than another scenario, this information can be maintained. The  
22 priority (e.g. a rating from between 1 to 100, with 100 being the highest priority),  
23 should be equal to a best guess as to the percentage of the time that users will step  
24 through the scenario, assuming they use the extension at all. Checkpoints provide  
25 a way to separate one scenario from another. For example, checkpoints designated

1 as “Offline” and “Shutdown” can be automatically added at the beginning and end  
2 of the scenarios, respectively, so there can be a differentiation between scenario  
3 runs in the log. Additionally, script authors can optionally use checkpoints mid-  
4 scenario to indicate a change in group priority, e.g. part of a scenario script could  
5 be labeled as an “On demand” feature and another part could be labeled as  
6 “Offline”.

7 Fig. 12 is a flow diagram that describes steps in a package manifest creation  
8 method in accordance with the described example. The method can be  
9 implemented in any suitable hardware, software, firmware or combination thereof.  
10 In the described example, various steps of the method are implemented by a  
11 manifest creation tool that is implemented in software.

12 Step 1200 provides a package manifest creation tool. This tool can be a  
13 software tool that resides on an extension author’s machine. Step 1202 receives  
14 information pertaining to an extension directory as a first input parameter. Step 1204  
15 determines whether there is any file group information or load dependency  
16 information provided by the extension author. If there is, step 1206 receives the  
17 information as an input parameter. Step 1208 determines whether there is any file  
18 usage statistic information. Such information can be provided, in one  
19 embodiment, through the use of scenario runs as described above. If such  
20 information is provided, then step 1210 receives the information as an input  
21 parameter. Step 1212 then uses all of the information provided as input  
22 parameters to generate a manifest automatically.

1                   Exemplary File Ordering Heuristics Based on File Usage Statistics

2                   Fig. 13 is a flow diagram that describes but one exemplary file ordering or  
3                   sorting heuristic in accordance with the described embodiment. It is to be  
4                   understood that this specific example constitutes but one way of ordering files for  
5                   download. Accordingly, other heuristics can be used without departing from the  
6                   spirit and scope of the claimed subject matter.

7                   Step 1300 sorts files by file group. Recall that in the illustrated example  
8                   above, files can be grouped in one of four possible groups: Required, Offline, On  
9                   Demand and Online Only. A file's group is determined first by the manifest, and,  
10                   if it does not provide any group information, then by the highest priority group  
11                   that it uses, according to checkpoint information in the log. Files in the  
12                   “Required” set should not be considered because their order is already known. If  
13                   no group information is included about a file, then an assumption is made that the  
14                   EDF and all DLLs are “Required” files and all other files in the directory are  
15                   “Offline”.

16                   Consider, for example, the following initial file usage information for three  
17                   different scenarios:

18                   Scenario 1 file usage: 1) FileA.gif, 2) FileB.xml, 3) FileE.dll

19                   Scenario 2 file usage: 1) FileC.xml, 2) FileA.gif

20                   Scenario 3 file usage: 1) File D.js, 2) FileA.gif

21                   Scenario 1 = priority 80

22                   Scenario 2 = priority 80

23                   Scenario 3 = priority 40

24                   In this example, there are three scenarios that have files associated with  
25                   them. Each of the scenarios has a priority with which it is associated. The files

are first sorted by group (step 1300). Recall that in this ordering heuristic, DLLs are “Required” and all other files are considered “Offline”. This provides the following sorted files:

## Required files

## FileE

## Offline files

FileA, FileB, FileC, File D

Step 1302 sorts files based on scenario priorities (from highest to lowest). Higher priority files are ordered so that they are downloaded first. This step provides the following sorted files:

### Required files

---

FileE

## Offline files

Priority 80 group: files used by Scenarios 1 & 2 = File A, File B, and File C

Priority 40 group: files used by Scenario 3 (that are not already listed) = File D.

Step 1304 then sorts the files by file usage order within a scenario run. For each priority grouping with more than one file, the files are sorted according to the average order in which they were downloaded within scenarios of their labeled priority. Scenarios with a smaller average usage order will be downloaded earlier. Ties are broken based on the order in which the scenarios appear in the input file. As an example, consider the following:

1 File A: average order = (Scenario 1 order + Scenario 2 order)/2 =  
1 (1+2)/2 = 1.5.

2 File B: average order = (Scenario 1 order)/1 = (2)/1 = 2.

3 File C: average order = (Scenario 2 order)/1 = (1)/1 = 1.

4 Here, file A got used first by scenario 1 and second by scenario 2 for an  
5 average of 1.5, and so one. File C has the smallest order number so, of the Offline  
6 files, it is sent first. The final file order is shown below:

7

8 Required files

9 FileE

10 Offline files

11 FileC, FileA, FileB, File D

12 **Code, Components and “bits”**

13 The following files and resources can be, but need not be included with an  
14 extension. This list is not exclusive, as other resources can certainly be  
15 incorporated into an extension.

16

17 • Customized UI and keyboard shortcuts

18 • Components and Behaviors

19 • XML browsing and editing components (including XSL and  
business logic objects)

20 • Static pages or other resources

21 • Third-party defined custom content

22 Users install extensions by navigating to a network site for the extension.  
23 In an Internet implementation, the user navigates to an appropriate URL for the  
extension. Hosting administrators can also “push” extensions so that users can

1 automatically receive them by adding an entry into the appropriate users'  
2 "Preference" settings.

3

4 **Platform Set Up and Extension Installation**

5 Fig. 14 is a flow diagram that describes exemplary steps in a set up and  
6 extension installation process in accordance with the described embodiment. This  
7 example describes an Internet-based example. In the illustrated example, various  
8 extensions are maintained at or accessible through various Internet sites. The  
9 extensions are deliverable via the Internet to a client. It will be appreciated that  
10 the illustrated division of computers may not necessarily exist at all. For example,  
11 all of the functionality embodied by the computers may reside on one machine, or  
12 the extension may be local, or the platform and the extensions may be on the same  
13 machine, etc.

14 The flow diagram, in this example, is illustrated in connection with three  
15 separate "zones", one of which represents a client, one of which represents a  
16 "platform" Internet server, and one of which represents a third party Internet  
17 server. The acts that are described in connection with the different zones are  
18 performed, in this example, by the entities assigned to the zone. In some  
19 configurations, one or more of these zones may overlap. For instance, the  
20 platform server may be the same device as the extension server.

21 Step 1400 navigates a user to a particular Internet site that is associated  
22 with the software platform that is to be utilized as the foundation for extension  
23 installation described below. In step 1402, the user clicks an "install" button that  
24 sends a message to the software platform server that indicates that a user wishes to  
25 install the software platform. This step can be an optional step. Steps 1404 and

1 1405 then download the software associated with the software platform to the  
2 client. In the illustrated example, step 1404 downloads the package file for the  
3 single navigable window application and based on that file's contents step 1405  
4 downloads other components and files to the user's computer. Step 1406 installs  
5 the software code on the client machine and can create local directories for the  
6 application cache, a local store and preferences. It will be appreciated, however,  
7 that local directories or preferences are not necessarily needed. Step 1408  
8 launches the software platform.

9 The steps described immediately above constitute steps associated with an  
10 initial set up in which the software code for the single navigable window  
11 application is delivered to and installed on a client machine. The steps described  
12 immediately below are associated with extension installation.

13 Step 1410 uses a link that is associated with an extension to access the  
14 extension. This step can be implemented by a user navigating their browser to a  
15 particular Internet site through which one or more extensions can be accessed.  
16 Alternately, a reference to the link can be placed in the user's preferences or the  
17 preferences of a computing group with which the user is associated (e.g. the  
18 system administrator can place a reference in a group's preferences). The link can  
19 advantageously be associated with a third party Internet server or Internet site.  
20 Step 1412 downloads extension files according to the PKG associated with an  
21 EDF. The files are delivered to the client and step 1414 places the extension files  
22 in a local store as specified by the PKG specification. At this point, an extension  
23 is installed and the user can utilize the functionality provided by the extension.  
24 Step 1416 determines whether extension updates are available. This can be done  
25 by periodically polling an extension catalog (discussed in an "Extension Catalog"

1 section below) to ascertain whether there are any updates to various extensions.  
2 Alternately, notifications might be automatically sent to the client so that the client  
3 is aware of updates or any other method might be used to determine if updates are  
4 available. If there are updates available, step 1418 branches to step 1412 which  
5 downloads the extension files associated with the update and installs them on the  
6 client.

7

### 8 Developing Extensions

9 Developing extensions for the software platform is a fairly straight-forward  
10 process. A developer develops the extension content using a tool such as Notepad  
11 or other tools such as Visual Studio. The extension is then described in an EDF  
12 and PKG and the PKG is digitally-signed and then optionally compressed. The  
13 extension can then be hosted on a particular network server.

14 Fig. 15 is a flow diagram that describes steps in an extension development  
15 method in accordance with the described embodiment. One or more of these steps  
16 can be performed by a software developer or organization that creates a particular  
17 extension. Some of the steps are implemented in software. Step 1500 develops an  
18 extension. Any suitable tools can be used to develop the extension. Step 1502  
19 then creates an extension definition file (EDF) for the extension. The EDF, in this  
20 example, is defined using XML as discussed above. Other formats can, of course,  
21 be used to describe the EDF. Step 1504 creates a package manifest (PKG) for the  
22 extension. The PKG, in this example, is defined using XML as discussed above.  
23 Step 1506 then hosts the EDF and PKG on a network server such as an Internet  
24 server. Additionally, the associated extension files that are described in the PKG  
25 can also be hosted by the network or Internet server (step 1508). Having

1       accomplished the above, users can now navigate to an EDF directly (using, for  
2       example, the associated URL or some other network address), which then installs  
3       the extension by caching any required files locally and placing a reference to the  
4       extension in the user's preferences.

5       Specifically, step 1510 delivers the EDF and PKG files to a client. This  
6       step can be implemented by a user navigating to a particular Internet site where the  
7       appropriate files are hosted and then downloading the files. Step 1512 delivers the  
8       extension files that are associated with the EDF and PKG files to the client,  
9       whereupon they can be installed and used.

10

11       **Extension Catalog**

12       One optimization, discussed briefly in connection with Fig. 14, is an  
13       extension or EDF catalog which provides an additional level of indirection to the  
14       EDF. An EDF catalog allows organizations to group extensions and provides a  
15       single place to determine when an extension changes. The desired extension can  
16       be automatically selected from the catalog by the software platform based upon  
17       the user's settings. The catalog can be queried to determine which extension is  
18       most appropriate for the user.

19       In the described embodiment, a catalog is an XML file that contains  
20       mappings from extension URNs to one or more package URNs based upon  
21       language, version or other attributes. Catalogs can, however, be defined using any  
22       suitable format. Catalogs can provide:

23

24       The ability for a hosting organization to update version information  
25       for one or more hosted extensions in a single place

1                   Optional automatic indirection to correct version based upon user's  
 2                   settings. For instance, a catalog may list several versions of an  
 3                   extension for different languages. The catalog file can be processed  
 4                   to find the version of the extension that matches the user's language  
 5                   settings.

6

- 7                   • Optional automatic upgrade to new versions of extensions as they  
 8                   become available

9

10                  Like EDFs, catalogs can be compressed and digitally signed to prevent  
 11                  tampering. By subscribing to a catalog in order to detect changes for one or more  
 12                  hosted extensions, the number of server pings required from the client (or  
 13                  notifications to the client) in order to discover extension updates can be reduced.

14

15                  Fig. 16 shows an exemplary XML catalog structure. Entries in the catalog  
 16                  can be structured as follows:

12 <b>Attribute</b>	13 <b>Type</b>	14 <b>Required</b>	15 <b>Description</b>
16                  extensionURN	17                  uri	18                  Y	19                  Identifier for an extension. There may be more than one entry for a given extension urn in a single catalog representing different versions, languages, etc.
20                  name	21                  String	22                  N	23                  Friendly name for an extension.
24                  packageURN	25                  uri		Required urn for the package. Package urn corresponds to a discrete set of bits. It is different from the extension urn: For each user, the extension urn (name) corresponds to a specific set of files based upon language preferences and version. This means that, for shared machines, different users may have different extensionURN to packageURN maps based

			upon their preferences.	
1	packageURL	uri	Y	url of the digitally signed compressed file containing the PKG file is required.
2	language	String	N	Language is an optional language specifier.
3	version	String	N	Version is an optional version specifier
4	defaultLanguage	String	N	DefaultLanguage is an optional attribute specifying the default language package. For a given version of an extension, there should be only one entry with DefaultLanguage attribute.
5	defaultVersion	String	N	DefaultVersion is an optional attribute specifying default version for an extension. For a given extension urn and language attribute there should be only one entry with DefaultVersion attribute.
6				
7				
8				
9				
10				
11				
12				
13				
14				
15				
16				
17				
18				
19				
20				
21				
22				
23				
24				
25				

In this particular example:

- The default language of the netdocs-planner is the English version.
- The default English version is 1.1. The default French version is 1.0. If there is no version available in the user's specified language on the platform, they will get the English version 1.1 by default.
- The English version of netdocs-planner has been upgraded from V1 to V1.1.
- There is also a French version. The extension URN is the same as the English version. There is no 1.1 release for French yet, so 1.0 is the current version for French speaking users.
- A query against the catalog returns only the rows where language matches the user's language preferences. The query would also return all rows where language is the user's language or default = 'yes' and throw out duplicates for the same name.

## 1                   Architecture

2                   In the described embodiment, one of the aspects that provide desirable  
3                   utility is the extensibility of the software platform. That is, third and fourth party  
4                   developers are free to develop their own extensions which can then be used within  
5                   the framework of the software platform. The extensions are integrated directly  
6                   into the software so that the platform's functionality is modified by the extensions.  
7                   Recall that in order to provide an extension, the developer simply authors the  
8                   extension, describes their extension in an EDF and PKG, and then hosts the EDF,  
9                   PKG, and associated files on a network server.

10                  The EDF, as pointed out above, can be defined in an XML schema that  
11                  includes a root node (i.e. the "extension" tag) and one or more child nodes. In this  
12                  particular example, the child nodes generally correspond to the individual  
13                  extension feature types that are desired for incorporation into the software  
14                  platform. For example, recall that Tables 1-3 above describe various exemplary  
15                  predefined feature types that can be added through an extension using the  
16                  predefined XML schema.

17                  Consider now a developer who wants to add two menus and a toolbar to the  
18                  software platform. The menus and toolbar might be associated with a retail store  
19                  that maintains a Web site for its customers. The retail store might want a customer  
20                  who visits its Web site to be presented with a UI that is unique to the retail store  
21                  and provides services that are specifically tailored to the store. To do this, the  
22                  developer develops two different menus, one of which might be associated with  
23                  displaying the most recent specials, and other of which might be associated with  
24                  providing a search mechanism through which the user can search for specific  
25                  products. The toolbar might contain specific buttons that are unique to the retail

1 store. A simplified EDF called “*retail.edf*” for the retail store’s extension is  
2 shown directly below:

```
3  
4 <edf:extension name= “retail extension” urn= “extension.retail.com”>  
5     <edf:menus>  
6         <edf:menu url= “url1.htm”/>  
7         <edf:menu url= “url2.htm”/>  
8     </edf:menus>  
9     <edf:toolbars>  
10        <edf:toolbar url= “url3.htm”/>  
11    </edf:toolbars>  
12 </edf:/extension>
```

13  
14 Here, the outer “extension” tag designates this XML file as an extension.  
15 The inner “menus” and “toolbars” tags are top level tags that designate that the  
16 information between these tags pertains respectively to menus and toolbars that  
17 correspond to the extensions that the developer has added. The boldface “menu”  
18 and “toolbar” tags describe data pertaining to the actual extension and contain a  
19 URL that is associated with each extension as described above. The EDF above  
20 logically describes the extensions that are being provided as including two menus  
21 and one tool bar.

22 Consider also that the above EDF can constitute but one of many EDFs that  
23 are loaded into the system. Each EDF can contain one or more top level tags, each  
24 of which is associated with one or more specific extensions that are to be added to  
25 the software platform.

26 Fig. 17 is a block diagram of an exemplary software architecture that is  
27 configured to process multiple different EDFs so that the software components  
28 that are responsible for incorporating each particular extension into the software

platform receive the appropriate information that is specific to their extension. This example is specific to the XML implementation that is discussed throughout this document. It is to be understood that other architectures, similar in functionality to the one discussed below, can be used in other implementations without departing from the spirit and scope of the claimed subject matter.

Utility objects, herein referred to as attachment points, are used to process the information from the multiple EDFs. An attachment point is simply a collection of objects that fire events to registered listeners as objects are added to or removed from the collection. Many types of attachment points can be created, but all take data from a source (often another attachment point), process it (either dynamically or statically), and expose the results of their processing. Some of the simplest attachment points include:

- An XML attachment point, which loads an XML file and exposes the top-level nodes of the XML as objects in its collection.
- A filter attachment point, that connects to another attachment point and exposes only those objects from it that meet some criteria.
- A merge attachment point, that connects to one or more other attachment points and exposes all of their objects as one, merged collection of objects.

In the illustrated example, the architecture includes a collection of one or more attachment points, including a funnel structure known as an EDFHub 1700, an attachment point manager 1702, and multiple attachment managers 1704. The EDFHub 1700 receives all of the EDFs and merges them together and exposes them as a single list. Other individual attachment points provide mechanisms that manipulate (including filter, merge and expand) the single list that is exposed by the EDFHub 1700. Whenever a new extension or EDF is added to or removed from the EDFHub, the various attachment points will see to it that the appropriate attachment manager(s) is notified. This is done by firing events to the appropriate

1 attachment managers. The attachment point manager 1702 creates, destroys and  
2 manages the various attachment points in the system and allows easy reuse of  
3 identical attachment points.

4 For each top level tag (i.e. “menus” and “toolbars” tags), there is a  
5 corresponding attachment manager 1704 that uses data provided by the attachment  
6 points to incorporate a particular type of feature within the software platform.  
7 Each attachment manager requests a set of attachment points from the attachment  
8 point manager 1702. These manipulate the data exposed by the EDFHub 1700. In  
9 the illustrated example, the attachment points can be requested as a predicate chain  
10 that the attachment point manager uses to create and build a set of attachment  
11 points that operate on the data exposed by the EDFHub 1700.

12 Fig. 18 is a flow diagram that describes steps in a method in accordance  
13 with the described embodiment. The method is implemented in software and, in  
14 this example, by the software components of Fig. 17.

15 Step 1800 receives multiple EDFs. These files can be received in any  
16 suitable way. For example, a user can specify in their preferences particular  
17 extensions that they desire to be loaded when they are online. Alternately, a user  
18 might navigate using a link to a particular Internet site that recognizes that the user  
19 is executing a software platform that is configured to dynamically add the  
20 extensions. The EDFs, in this example, are funneled into the EDFHub 1700 which  
21 uses attachment points to combine the EDFs (step 1802). In this example, the  
22 EDFs are defined as XML files and the nodes are combined into one single XML  
23 list. Step 1804 exposes the combined EDFs. In this particular example, the EDFs  
24 are combined into a single XML list that is exposed to other various attachment  
25 points which then further manipulate that data (step 1806). One goal of the

1 attachment points is to avoid having the attachment managers 1704 re-query the  
2 whole system every time an extension gets added or removed from the system.  
3 Thus, if an extension gets added or removed, the attachment points ensure that  
4 only the appropriate attachment manager 1704 is notified of the specific additions  
5 or deletions of the extension. For example, if an EDF indicates the addition of a  
6 menu, then only the attachment manager associated with menus is notified.  
7 Accordingly, step 1808 notifies the appropriate attachment manager of any new  
8 data that matches the attachment managers requirements.

9

## 10 **Attachment Points and the Attachment Point Manager**

11 Attachment points are objects that expose collections of ordered objects and  
12 fire notifications when new objects are inserted or deleted. In the exemplary  
13 system, the objects are XML nodes, but they could be any type of object.  
14 Although there are many types of attachment points, they all follow a similar  
15 process:

- 16 1) Initially attach to one or more data sources. These could be files  
or, commonly, other attachment points.
- 17 2) Process the data based on some logic. Usually the logic is quite  
simple and could involve something like filtering the objects  
based on some criteria.
- 18 3) Expose the results of the processing step 2 in a new collection of  
objects.
- 19 4) Fire events to indicate how the exposed collection of objects  
changed (OnInserted(index, count) or OnRemoved(index,  
count)).
- 20 5) Optionally, continue to listen for changes in the data sources and  
repeat step 2-4 when changes occur.

23 Alone, each attachment point is quite simple, but when different types of  
24 attachment points are combined to form “chains”, where one attachment point  
25 processes data from a second attachment point, the processing can be quite

1 powerful. This is especially true if the attachment points only process the changed  
2 data in step 2, since they only do a small amount of simple work at any one time.  
3 In the exemplary system, this incremental processing means that the whole system  
4 does not have to be re-queried when a new extension is installed or when an  
5 existing extension is removed. Additionally, each attachment manager in the  
6 exemplary system uses a specific chain of attachment points and so is only  
7 informed of changes that impact its area of extensibility.

8 The attachment point manager performs two important functions when  
9 building chains of attachment points. First, it allows the chains to be described as  
10 predicate strings. The attachment point manager interprets these strings and builds  
11 the required chain of attachment points. Second, it allows the reuse of identical  
12 attachment points, which makes the system more efficient. As the attachment  
13 point manager creates each chain of attachment points, it tracks which predicate  
14 string corresponds to which attachment point. If that predicate string is later  
15 requested again, it simply reuses the existing attachment point without creating a  
16 new one.

17 As an example, consider that the attachment manager associated with  
18 menus has requested the following predicate chain of attachment points that  
19 utilizes the *retail.edf* file above: (Note: this example does not assume the presence  
20 of an EDFHub attachment point).

21  
22 *Explode(Filter("menus", Explode (URL ("retail.edf"))))*

23  
24 This string represents all of the menus in the *retail.edf* file. The XML file  
25 located at *retail.edf* is loaded by the URL attachment point which exposes the root

1 node of the XML file as the only object in its collection. The inner Explode  
2 attachment point uses the URL attachment point as its data source and exposes all  
3 of the children of the objects in that source collection. In this case, the children of  
4 the root node are the top level XML tags “menu” and “toolbars”. The Filter  
5 attachment point uses the Explode attachment point as its data source and filters  
6 the exposed objects looking only for the nodes that are “menus”. The outer  
7 Explode attachment point uses the Filter attachment point as its data source and  
8 exposes all of the children of the filtered menu nodes to provide a list containing  
9 the two menus that are added by the extension. Since this particular XML file  
10 contained menus that were identified by the attachment points associated with  
11 menu attachment manager, that attachment manager is then notified that two  
12 menus have been added by an extension.

13 This process is diagrammatically illustrated in Fig. 19 which shows  
14 attachment points 1900, 1902, 1904, and 1906. Each attachment point exposes a  
15 list of XML nodes. URL attachment point 1900 takes an input (a URL to an XML  
16 file—e.g. *retail.edf*) and exposes a list of XML nodes. This list contains only the  
17 root node “<edf:extension>”. Explode attachment point 1902 takes as an input  
18 attachment point 1900 and exposes a list of XML nodes which are children of  
19 source XML nodes. In this example, the list of XML nodes exposed by  
20 attachment point 1902 are the “<menus>” nodes and the “<toolbars>” nodes. The  
21 filter attachment point 1904 takes attachment point 1902 as an input and filters on  
22 “menus.” It then exposes an XML list having only “<menus>” nodes in it. The  
23 explode attachment point 1906 takes attachment point 1904 as an input and  
24 exposes a list with the XML nodes that are contained in the “<menus>” nodes—  
25 here both of the “<menu>” nodes.

1 Consider additionally that the toolbar attachment manager would request a  
2 predicate chain of attachment points which would also use URL attachment point,  
3 an Explode attachment point and a filter attachment point 1904 that filters on  
4 “toolbars”. Thus, the corresponding explode attachment point 1906 would expose  
5 an XML list containing only the “<toolbar>” node. But, the attachment point  
6 manager would detect the commonality of the URL attachment point and the inner  
7 Explode attachment point, so it would reuse the same attachment points it created  
8 for the menu attachment manager. The Filter attachment points used by the  
9 toolbar attachment manager and the menu attachment manager would use the  
10 same Explode attachment point as their data sources but would expose different  
11 collections of nodes, because they were filtering based on different criteria.

12 Consider Fig. 20 which incorporates an EDFHub attachment point 2000.  
13 This attachment point receives all of the EDFs and, as discussed above, combines  
14 them into a single XML list. The EDFHub then exposes the root nodes of all of  
15 the EDFs. The explode attachment point 2002 then exposes an XML list that  
16 contains all of the top level nodes for all of the EDFs. As an example, there may  
17 be multiple EDFs that each contain top level menu nodes, toolbar nodes,  
18 accelerator nodes and the like. Explode attachment point 2002 exposes an XML  
19 list that contains all of these top level nodes for all of the EDFs. Filter attachment  
20 point 2004 can then filter the XML list exposed by the explode attachment point  
21 2002 in accordance with any suitable parameters (i.e. filter on menu nodes, tool  
22 bar nodes, accelerator nodes and the like). The final explode attachment point  
23 2006 then exposes a list of the individual children nodes of the list exposed by the  
24 filter attachment point 2004. This list describes all of the specific features (of the  
25 particular type that were filtered) that have been added by all of the EDFs.

1           The table below lists a number of different attachment points that can be  
2           utilized in accordance with this described embodiment but many more can easily  
3           be created.

5 <b>Attachment Point</b>	6 <b>Purpose</b>
7           URL	8           Loads the URL and exposes the top 9           level XML node as a member of the 10          collection
11          Context	12          For every member, it gets the 13          “expression” attribute and binds to it. If 14          the expression evaluates to true, then 15          the member is exposed.
16          EDF	17          Same as the URL AP, but also exposes a 18          fabricated member with data to create 19          an APP based on the URL and URN 20          (which exists in the XML DOM).
21          Merge	22          Takes zero or more Attachment Points 23          (of any type) and merges them together. 24          The order and continuity of the original 25          collections will be maintained.
26          Filter	27          Monitors a single Attachment Point and 28          only exposes those nodes that match the 29          specified name. The order of the 30          original collection will be maintained.
31          Duplicate	32          Monitors a single Attachment Point and 33          filters out any duplicates. A duplicate is 34          defined to be having the same URN 35          attribute. If no URN attribute is present 36          then the node is exposed. Order of the 37          original collection will be maintained.
38          Explode	39          Monitors a single Attachment Point and 40          for every member exposes the children 41          of that member as its members. The 42          order of the original collection will be 43          maintained as well as the order of the 44          children within the nodes.
45          Link	46          Monitors a single Attachment Point and 47          for every member looks for a URL 48          attribute and creates a URL AP and

1	merges it into itself. If the optional 2 include Content is set to true, it will 3 merge the original collection in as well.	
4	Order 5 6	Monitors a single Attachment Point. 7 For every member, it gets three 8 attributes: id, before and after. Based 9 on this information, it reorders the 10 members as specified. If no ordering 11 information is supplied, the order of the 12 original collection will be maintained.
13	EDFHub 14	This Attachment Point is the central 15 merge point that represents all the EDF 16 points.

## 9 Conclusion

10 The embodiments described above provide a platform solution that  
11 provides for customization and extensibility through a consistent and logical  
12 extensibility mechanism and object model that can be easily understood by third  
13 party developers. Internet-based downloads can be accomplished without a great  
14 deal of user intervention and without manipulating any user persisted settings.  
15 Extensions can be provided to a software platform or application program  
16 dynamically based upon the user's computing context.

17 Although the invention has been described in language specific to structural  
18 features and/or methodological steps, it is to be understood that the invention  
19 defined in the appended claims is not necessarily limited to the specific features or  
20 steps described. Rather, the specific features and steps are disclosed as preferred  
21 forms of implementing the claimed invention.